

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (previously presented) A method for forwarding packet-based traffic through a network node, comprising:
 - receiving traffic type bandwidth limitations from a customer;
 - dedicating a group of queues in a network node to the customer;
 - translating the traffic type bandwidth limitations, which were received from the customer, to queue-specific bandwidth limitations;
 - performing queue-specific rate shaping on the customer's traffic according to the queue-specific bandwidth limitations respectively associated with the queues; and
 - performing group-specific rate shaping on the customer's traffic according to a group-specific bandwidth limitation associated with the group of queues.
2. (original) The method of claim 1, further comprising associating queues from said group of queues with different types of traffic that are to be received from the customer.
3. (canceled)
4. (original) The method of claim 1, further comprising associating said group of queues with a group rate shaper that performs said group-specific rate shaping on the customer's traffic on an aggregate basis.
5. (original) The method of claim 1, further comprising prioritizing the queues of said group of queues.

6. (original) The method of claim 5, further comprising:
distributing said portion of excess unused bandwidth among the group of queues on a priority basis according to said prioritizing.
7. (original) The method of claim 1, further comprising:
scheduling packets for forwarding from one or more of said queues in said group of queues, wherein bandwidth consumed by the packets from each of the queues is less than or equal to respective queue-specific bandwidth limitations for the queues;
identifying excess unused bandwidth when the consumed bandwidth is less than said group-specific bandwidth limitation; and
distributing a portion of the excess unused bandwidth to a first queue of the group of queues, wherein the sum of the consumed bandwidth and the portion of the excess unused bandwidth is less than or equal to a group-specific bandwidth limitation for the group.
8. (previously presented) A network node for forwarding packet-based traffic, comprising:
a plurality of queues;
a plurality of queue-specific rate shapers respectively associated with the plurality of queues;
a plurality of group-specific rate shapers configured to be associated with groups of the plurality of queues;
a group establishment module configured to dedicate a group of said queues to a customer and to associate one of said group-specific rate shapers with said group of queues that is dedicated to said customer; and
a scheduler configured to:
schedule, in a first round, packets enqueued in the plurality of queues according to the respective plurality of queue-specific rate shapers and the respective group-specific rate shapers; and
schedule, in a second round, packets enqueued in the plurality of queues according to the respective group-specific rate shapers;

wherein said scheduler is further configured to:

schedule, in subrounds of the first round, packets enqueued in the plurality of queues according to a priority respectively associated with each of the queues and schedule, in subrounds of the second round, packets enqueued in the plurality of queues according to the priority respectively associated with each of the queues;

wherein the scheduler comprises an individual queue enablement vector for each queue, a group enablement vector for the group of queues, and a result vector for each queue;

wherein the individual queue enablement vector indicates which queues are enabled, with a queue being enabled if the queue has not consumed its allocated queue-specific bandwidth;

wherein the group enablement vector indicates whether the group is enabled with the group being enabled as long as all of the allocated group-specific bandwidth has not been consumed; and

wherein the result vector indicates which queues are enabled for sending packets, wherein in the first round a result vector for a queue indicates a queue is enabled only when both the individual queue enablement vector and the group vector indicate that the queue is enabled and in the second round a result vector for a queue indicates a queue is enabled as long as the group vector indicates that the group is enabled.

9. (original) The device of claim 8, further comprising:

a scheduler, coupled to the plurality of queue-specific rate shapers and the plurality of group-specific rate shapers, configured to schedule packets enqueued in the plurality of queues according to the respective plurality of queue-specific rate shapers, wherein the queue-specific rate shaper respectively associated with each queue is associated with a priority, and wherein the scheduler schedules according to the associated priority.

10. (original) The device of claim 9, wherein said scheduler is further configured to:

scheduling packets for forwarding from a first one or more queues of said plurality of queues, wherein bandwidth consumed by the packets from each of the first one or

more queues is less than or equal to respective queue-specific bandwidth limitations for the first one or more queues;

identifying excess unused bandwidth when the consumed bandwidth is less than a group-specific bandwidth limitation, wherein a sum of the consumed bandwidth and the excess unused bandwidth approximately equals the group-specific bandwidth limitation; and

scheduling packets for forwarding from a second one or more queues of said plurality of queues using the excess unused bandwidth.

11. (canceled)

12. (canceled)

13. (canceled)

14. (original) The device of claim 8, further comprising:

a plurality of pipes, wherein each pipe is associated with a group-specific rate shaper, and wherein each pipe of said plurality of pipes includes:

multiple traffic channels comprising one or more queues of the plurality of queues, wherein each traffic channel is associated with a queue-specific rate shaper.

15. (previously presented) A method for packet-based traffic forwarding, comprising:

establishing a customer-specific bandwidth limitation for a customer;

receiving traffic-type-specific bandwidth limitations from the customer;

dedicating multiple traffic channels to the customer;

associating the customer-specific bandwidth limitation to the traffic channels;

associating the traffic-type-specific bandwidth limitations with the traffic channels;

performing traffic-type-specific rate shaping according to the traffic-type-specific bandwidth limitations respectively associated with the traffic channels; and

performing customer-specific rate shaping according to the customer-specific bandwidth limitation associated with the traffic channels.

16. (original) The method of claim 15, further comprising:
prioritizing the traffic channels relative to one another.
17. (original) The method of claim 16, wherein said performing traffic-type-specific rate shaping consumes less bandwidth than said customer-specific bandwidth limitation, said method further comprising:
identifying excess unused bandwidth following the traffic-type-specific rate shaping;
and
distributing the excess unused bandwidth to a subset of the traffic channels in priority order according to said prioritizing.
18. (original) The method of claim 15, further comprising:
associating a traffic type with each traffic channel.
19. (original) The method of claim 18, further comprising:
adjusting the traffic-type-specific rate shaping according to traffic type-specific rate shaping customer preferences.
20. (previously presented) The method of claim 15, further comprising:
associating respective traffic-type-specific bandwidth limitations with each traffic channel, wherein the sum of the respective traffic-type-specific bandwidth limitations is less than or equal to the customer-specific bandwidth limitation.